Comparison of Global Precipitation Estimates across a Range of Temporal and Spatial Scales

Maria Gehne*

CIRES, University of Colorado, Boulder, Colorado

Thomas M. Hamill

National Oceanic and Atmospheric Administration, Boulder, Colorado

George N. Kiladis

National Oceanic and Atmospheric Administration, Boulder, Colorado

Kevin E. Trenberth

National Center for Atmospheric Research, Boulder, Colorado

**Corresponding author address:* CIRES, University of Colorado, Boulder, Colorado.

E-mail: maria.gehne@noaa.gov

ABSTRACT

Characteristics of precipitation estimates for rate and amount from three global High-resolution precipitation products (HRPPs), four global Climate Data Records, and four reanalyses are compared here. All data sets considered have at least daily temporal resolution. Estimates of global precipitation differ widely from one product to the next, with some differences likely due to differing goals in producing the estimates. HRPPs are intended to produce the best instantaneous precipitation estimate locally. Climate data records of precipitation emphasize homogeneity over instantaneous accuracy. Precipitation estimates from global reanalyses are dynamically consistent with the large scale circulation but tend to compare poorly to rain gauge estimates as they are forecast by the reanalysis system and precipitation is not assimilated. As expected, variance and the average spread among data sets are highest where the means are large. Regionally, differences in the means and variances are as large as the means and variances respectively. Temporal correlation, rain rate and rain amount distributions, and biases in time evolution are explored using temporal and spatial averaging. It is shown that differences on annual time scales and continental regions are around 0.8mm/d, which correspond to 23W m^{-2}. These wide variations in the estimates, even for global averages, highlight the need for better constrained precipitation products in the future.

1. Introduction

Gridded estimates of daily (or higher frequency) global precipitation are becoming more and more necessary for applications such as model validation, input for land-surface models, or extreme-event characterization. Detailed knowledge about current precipitation distributions is also necessary to quantify changes in precipitation estimated by global-warming scenarios, which tend to be described as changes in the mean and tails of the distribution. On monthly scales global precipitation estimates have been used to estimate the global water cycle (**?**), study the co-variability of precipitation and surface temperature (**?**), and to assess the imbalance between global precipitation and evaporation (**?**). All of these applications assume that an accurate or at least adequate estimate of these distributions is obtainable. For many other applications, higher temporal (sub-monthly) and spatial resolution is needed. Extreme precipitation events are usually highly localized in space and time, involving temporal scales on the order of minutes to a few hours and several kilometers, especially in summer over land. To resolve the more extreme precipitation intensity events data on ten minute intervals thus might be needed (**?**). To accurately identify the mean diurnal cycle, hourly time steps are desirable. The highest resolutions of current global precipitation estimates are 3 hours and 0.25°. This is marginally adequate to resolve the diurnal cycle and mesoscale systems but is still too coarse to resolve individual mesoscale storms. For model purposes, where time steps may be on the order of half an hour, hourly data are best. Hourly resolution sets a good compromise between what is meaningful for models and useful for extremes.

Gridded rain-gauge based analyses of precipitation are available over the global land areas, with the estimates assumed to be representative for a given area. Individual rain-gauge estimates of precipitation exist in many locations, but these are point estimates and apply only at the location they were collected. Large land areas on the globe are very sparsely covered with rain gauges, and ocean areas are not covered at all. In sparsely sampled areas, interpolation between rain gauge locations to obtain a gridded analysis may introduce errors. In addition, rain-gauge estimates are thought to underestimate very high rain rates due to under-catch in high-wind or snow conditions (e.g. **??**). Another issue is that precipitation measurements are usually reported only once or twice a day, which affects both rates and totals, because the longer the rain is left in the gauge the greater the potential for some of it to evaporate. To resolve the very high rain rates in thunderstorms, for example, temporal resolution of hours or even minutes is necessary. Overall, gauge-based analyses are likely to be quite accurate in data-dense areas and questionable in data-sparse areas. Other available options for global precipitation estimates, that provide higher spatial and temporal resolution, are based on satellite data. Global reanalyses offer another way to estimate global precipitation with the advantage that all variables are somewhat dynamically consistent. These estimates are also available over the oceans.

There are several important questions users of these data sets need to ask. The most important one is obviously, which of these estimates is closest to the truth? There is no clear answer to this, even among satellite precipitation data sets. The conclusion of several precipitation inter-comparison projects was that no one methodology is superior to the others (**?**). **?** showed for regional comparisons, that uncertainty in the ground validation data are larger than the passive microwave (PMW) algorithm bias. They also showed that the differences in estimated rain rates are mainly due to how the more intense rain rates are calculated and how strict the screen (precipitating and dry pixels) is. On monthly timescales for global analyses, **?** show that merged analysis products, using more than one satellite source and rain gauge adjusted, are superior to single source products. Without the adjustment to rain gauges, large biases exist over the southern Great Plains in the US for high resolution precipitation products (**?**). Even rain gauge only data sets have large differences; in the context of drought, using one or another data set can mean an observed increase or decrease in drought (**?**). The main conclusion from these studies is that there is no best product, there is only the most appropriate product for a certain purpose. Given that no one product is perfect for all circumstances, a question that may be more appropriate to ask by the user, and more likely to yield a useful answer, is, which of these should be used for a particular application? For example, studies at different locations and different seasons will likely benefit from using the product that has been shown to do well under those conditions. If the emphasis is on consistency of precipitation with circulation patterns, then reanalysis products combined with observed precipitation may be the best choice. In addition, several other issues are not addressed in these previous studies. Are there systematic biases among the high-resolution precipitation estimates on the global scale?

In all cases it is important for the user to know how the products differ in their precipitation estimates. In order to answer this question it is necessary to first quantify the differences among the data sets and the different estimation approaches. Are there biases that are particular to a certain approach to precipitation estimation? How do the distributions differ? And, given all the different estimates, is there a way to quantify the uncertainty associated with them?

The aim of this study then, is not to determine which data set is closest to the absolute truth since that is impossible, but rather to identify strengths and shortcomings of the data sets, and to provide some guidance as to which data sets are likely to perform better in certain situations. We are interested in global precipitation data sets with daily or higher resolution. Global products are consistent for all areas of the world. This consistency helps in comparing different precipitation regimes across the globe, as the differences are not related to different analysis algorithms. Daily or higher temporal resolution is better suited for estimating distributions, than monthly resolution.

Section 2 introduces the data sets used in this study. Section 3 has the details of the statistics used to compare the precipitation estimates and how the distributions are computed. Section 4 evaluates the statistics and distributions, mostly on the example of North America, but other continental regions are mentioned to highlight stark differences or similarities. Figures for all other continental regions are included in the supplementary material. Lastly, section 5 summarizes and discusses the implications of the results presented in this study.

2. Data Sets

The lowest native resolution of all precipitation data sets under consideration here is daily on a 1° grid. Therefore, all data sets were interpolated from their original grids to a grid with 1° spatial and daily temporal resolution using conservative averaging. This was done to facilitate comparison of distributions and variability, to ensure that the precipitation estimates are comparable and to minimize biases. As temporal averaging is done to daily resolution, differences in the diurnal cycle phase and amplitude will not be resolved; the resolved time scales that will be considered are daily to interannual. The seasonal cycle has a large effect on precipitation, which is why all analyses are performed for each month of the year separately.

Our criteria (global data, daily resolution) exclude several well established precipitation estimates from this study, for reasons related to either their temporal resolution or their regional coverage. These include PRISM (**?**), the North American regional reanalysis (**?**), stage IV radar data (**?**), and Asian Precipitation - Highly Resolved Observational Data Integration Towards Evaluation of Water Resources (APHRODITE, **?**), because they are regional products, and the Global Precipitation Climatology Centre (GPCC, **?**) , GPCP monthly estimates (**?**), CPC merged analysis of precipitation (CMAP, **?**) and CRU precipitation (**?**), because of their monthly resolution.

a. High-resolution precipitation products

High-resolution precipitation products (HRPPs) aim to provide the best instantaneous precipitation estimates at high spatial and temporal resolution. Commonly, high-resolution infrared (IR) brightness temperatures from geostationary satellites are related to precipitation rates using the more accurate passive microwave (PMW) estimates from the polar-orbiting satellites. How these measurements are related, how the IR is calibrated, and whether the monthly means are scaled to match monthly rain gauge analyses varies between algorithms and constitutes the main sources of differences between the estimates; see **?** for an overview and an in-depth description of the various techniques. In general, PMW gives a more accurate instantaneous estimate of precipitation than IR, because of the more direct observation of precipitation. But this accuracy deteriorates for longer time averages due to the lower sampling frequency of PMW. The combination of PMW and IR measurements includes the different errors inherent in each technique (**?**).

The Climate Prediction Center morphing method (CMORPH, **?**) estimates rainfall by combining IR and PMW measurements. High-quality PMW rainfall estimates are propagated (using linear interpolation in time) by motion vectors derived from high frequency IR imagery. CMORPH is available from 2003-2013 at 3-hourly intervals on a 0.25° grid from 60°S to 60°N.

The Tropical Rainfall Measuring Mission (TRMM) 3B42 product, provides 3-hourly precipitation estimates on a 0.25° grid between 50°S to 50°N and from 1998 to 2013. The microwave-calibrated IR rainfall estimates use the same monthly satellite-gauge analysis as the Global Precipitation Climatology Project (GPCP, see below) to match the monthly totals. TRMM was previously determined to have large relative errors at small precipitation rates, however time/area averaging significantly reduces the random error (**?**).

The Precipitation Estimation from Remotely Sensed Information using Artificial Neural Networks (PERSIANN) algorithm merges high-frequency IR images with low frequency rainfall estimates from the TRMM satellite using artificial neural networks (**??**). The precipitation estimates are based on IR from geostationary satellites, and PMW measurements are used to update the algorithm parameters. PERSIANN is available from 2001-2013 at 3-hourly intervals on a 0.25° grid from 50°S to 50°N.

b. Climate data records of precipitation

For climate data records homogeneity is emphasized over instantaneous accuracy. The Climate Prediction Center (CPC) rain-gauge (GAUGE) data set is based on quality-controlled station data from more than 30000 stations. These data are then interpolated to create analyzed fields of daily precipitation with bias correction for orographic effects (**?**). The global analysis is available daily on a 0.5° grid from 1979-2005 (**??**). The real-time version of the CPC gauge data set (GAUGERT) uses about 17000 stations and is available on the same grid at the same time resolution from 2005-2013.

Global Precipitation Climatology Project (GPCP, v1.2) daily, 1° precipitation estimates are computed based on the threshold-matched precipitation index (TMPI) (**?**). For the TMPI, IR temperatures are compared to a threshold, and all cold pixels are given the same conditional precipitation rate, with threshold and conditional precipitation rate set locally by month. The monthly means are normalized to match the monthly GPCP satellite-gauge precipitation estimate (**?**), which is based on satellite data and rain-gauge analysis from the Global Precipitation Climatology Centre (GPCC). The GPCC monthly rain gauge analysis is bias corrected to account for systematic errors due to wetting, evaporation, or aerodynamic effects (**?**).

One of the latest climate data records is the Precipitation Estimation from Remotely Sensed Information using Artificial Neural Networks - Climate Data Record (PERSICDR, v1r1, **?**). This is generated using the PERSIANN algorithm, and adjusted using the GPCP monthly product to match monthly precipitation rates on a 2.5° grid between the two products. In contrast to the HRPP PERSIANN, the PERSICDR model is pretrained on stage IV hourly precipitation data and the model parameters are then kept fixed for the full historical record of IR data. PERSICDR is available on a 0.25° grid between 50°S to 50°N and from 1983 to present day.

c. Reanalysis precipitation products

Another way to estimate global precipitation is through short-term forecasts provided by global reanalyses. The underlying models assimilate a wide variety of observations, but in general not precipitation measurements or analyses. Precipitation is usually provided by a prior short-range forecast, and this inherits the systematic errors of the forecast model. The advantage to reanalyses is that all variables are somewhat dynamically consistent. However, as precipitation data are not typically constrained by the analysis procedure, reanalyzed precipitation is highly model dependent (**?**). This is particularly true in the tropics and over continents during the summer, when convective precipitation dominates. This leads to the well-known problem with precipitation estimates from general circulation models (GCMs) of raining too frequently, with an over-abundance of light rainfall and too infrequent extreme precipitation (e.g. **??**). As global reanalyses are based on similar GCMs they tend to have the same short-comings in this respect. One exception is the North American Regional Reanalysis (**?**), which does assimilate precipitation. And there is evidence that assimilation of precipitation significantly improves precipitation estimates and the atmospheric moisture budget (**???**).

To facilitate comparison of reanalyses with the other precipitation estimates, the reanalyses must be be generated at as high resolution as the other estimates. Lower-resolution reanalyses previously have been found to have lower rain rates and a smaller range of resolved rain rates overall as compared to satellite or gauge based estimates, similarly to operational forecast models (**?**). We obtained similar results when applying our analysis to lower resolution reanalyses. Here we consider the most recent global reanalyses products which have a spatial resolution of smaller than 1°. They are the European Centre for Medium-Range Weather Forecasting (ECMWF) ERA-Interim reanalysis (ERAI **?**), the Modern-Era Retrospective Analysis for Research and Applications (MERRA **?**), the NCEP Climate Forecast System Reanalysis (CFSR **?**), and the Japanese 55-year Reanalysis (JRA55 **?**).

d. Caveat to independence of precipitation estimates

None of the above precipitation estimates is independent of all the others, for there is a large degree of overlap in the source data that goes into the different estimates (Table 1). PERSIANN and CMORPH are the only satellite product without any ground validation with gauge data. Both TRMM and GPCP use the same algorithm and the same monthly satellite-gauge analysis to constrain their monthly totals (**?**). The GAUGE and GAUGERT estimates are for non-overlapping time periods and use a different total number of stations, but the underlying algorithm is the same. Their statistics compare very well even though only about half the number of stations are available for the real-time product GAUGERT (17000 compared to 30000 for the retrospective GAUGE).

3. Methods

The methods used to evaluate the precipitation estimates include basic statistical quantities such as means and variances, and their differences among products at each grid point (Table 2). We also show the mean and variance differences as percentage of the mean and variance respectively to compare their relative sizes. In addition we consider temporal averages on time scales of a week, a month and a year. Spatial averages are always area averages, taking into account the change in grid area with latitude.

Frequency distributions of precipitation are highly skewed, with the smallest rain rates being the most frequent. In general this makes comparing different distributions difficult, because the tails tend to be under-sampled. One way to reduce the discrepancy between the number of samples in the lower rain rate bins and the higher rain rate bins is to use logarithmic bin sizes that increase with rain rate. Of course, in that case care needs to be taken when computing integrals. In addition to frequency distributions of precipitation rate we also compare rain amount by rain rate distributions. The integral under the curve is equal to the total precipitation amount. These distributions tend to be skewed towards lower precipitation rates with the largest amounts occurring at intermediate rain rates. For both types of distributions a logarithmic bin size is used. The number of bins is 100 with a constant logarithmic (to base 10) bin length. Setting the minimum bin to 10^{-4} and the maximum to 10, the bin length then comes out to $\triangle b = \left(\log_{10} 10 - \log_{10} 10^{-4}\right)/100 = 0.05$. The edges of the bins are computed according to $b_i = 10^{-4} 10^{i \triangle b}, i = 0, .., 100$, which results in increasing bin sizes with precipitation rate. Rain rates below the minimum (including zero rain rates) are counted in the lowest bin.

Global maps of the spread among precipitation data sets (Table 2) can be used to identify regions with more or less variability among the data sets. First the mean seasonal cycle is removed from each data set. The spread is then computed as the standard deviation among data sets at each grid point and time and averaged for each month of the year.

4. Results

The continental regions used in the analyses are defined as the land areas contained within the latitude-longitude areas given in Table **??**.

a. Annual cycle

A summary of the annual cycle is given in Figs. **??** and **??** in form of its amplitude and phase. Differences in the amplitude are large over equatorial Africa and South America, and the Indian Monsoon region. Over North America the amplitude of the annual cycle in the midwest of the Unites States ranges between $3-13\mathrm{mm\,d^{-1}}$. The phase is defined as the day of the year the annual cycle is maximized, and so does not take into account if a location has multiple maxima in precipitation during the year. This is potentially an issue in equatorial South America and Africa, although overall the timing of the annual maximum in precipitation is captured consistently among the estimates. Regions with large discrepancies in timing are northern Africa, parts of Australia (both regions where the annual cycle amplitude is very small), and the northwestern United States (**??**).

b. Differences in means and variances

Distinctive differences among data sets of large-scale patterns of means and variances can be identified. The climatological mean monthly precipitation for July is shown in Fig. **??**. Comparison of the mean monthly precipitation across data sets shows large variability (Fig. **??**b-d), especially in areas like the Intertropical convergence zone (ITCZ). Other regions with large differences in the means are continental areas in the summer hemisphere and the western boundary ocean current regions. Because of large spatial gradients in some regions, small variations in the location of climatological features like the ITCZ can lead to large local differences in mean precipitation.

Figures **??**c,d and **??**c,d show that GPCP mean precipitation exceeds mean precipitation from satellite-only products PERSIANN and CMORPH over the oceans, except in regions with intense convective precipitation. TRMM and GPCP match well over land, but TRMM commonly has higher means over tropical oceans and smaller means over midlatitude ocean areas (Fig. **??**b). The closest match is between GPCP and PERSICDR monthly means (Fig. **??**f), where any differences are below 0.075mm d^{-1}. Satellite-only products PERSIANN and CMORPH have higher means over summertime continental regions than the gauge corrected estimates. Over land the main bias for gauge corrected precipitation estimates is due to the bias in the rain gauge analysis used. This is visible in the differences between GPCP monthly means and GAUGE monthly means (Figs. **??**e and **??**e), where the rain gauge analysis that contributes to GPCP is bias corrected for losses due to wetting, evaporation, or aerodynamic effects, and the CPC GAUGE analysis is corrected for orographic effects. Comparing the July estimates to January it becomes clear that CMORPH and PERSIANN tend to underestimate winter precipitation over continents and overestimate summer precipitation when compared to GPCP. GAUGE estimates are biased low compared to GPCP, and TRMM exceeds GPCP in regions of vigorous convection.

Percentage differences of the monthly means (Fig. **??**) show clearly that the differences in the means are often as large as the means. This is especially true in areas with small mean values like the subtropical dry zones, where small differences translate into large percentage differences. Depending on the data set under consideration, this can also be the case in regions with large mean precipitation and large variability like the continental US in the summer and the edge of the ITCZ (e.g. GPCP and CMORPH (Fig. **??**c)).

Monthly mean daily precipitation variance is large where mean precipitation is large (Figs. **??**a and **??**a). The largest variances are in areas with highly variable convective precipitation such as the ITCZ, the Indian Ocean, and the Indian Monsoon region. CMORPH has larger variance than all data sets except TRMM (Fig. **??**b,c), and differences in variances are as large as the variance for most areas of the globe (not shown). This holds even for areas with large variability, like the ITCZ. That magnitudes of spread and mean should correlate is to be expected for a positive definite quantity like precipitation, the magnitude of the difference in variance among data sets however is notable. Both rain gauge data sets show smaller variance than GPCP (Fig. **??**e). This is likely related to the fact that under catch for rain gauges tends to be more of an issue at higher rain rates, thus decreasing the variance. PERSICDR variance is smaller than GPCP variance over land, but exceeds GPCP variance over the ocean. Note, however, that differences in variance are smaller for PERSICDR and GPCP than for any other data set Fig. **??**f and **??**f). While CMORPH has the highest variance for most regions, Figs. **??**c and **??**c show that GPCP variance is higher in winter hemisphere. This issue will be discussed more in the following sections.

c. Time Series

Time series at the continental scale are shown for North America, where there is a relatively dense observing network and so the potential for constraining estimates is high. Time series averaged over North America are also a good example in that they illustrate many of the issues also observed in other regions. Other regions (Table **??**) are mentioned where results are notable, but results are not shown. Figures for all other regions are included in the supplementary material. Table **??** also includes the amplitude and phase of the mean seasonal cycle averaged over each continental region. The minimum and maximum amplitude estimated by the different products in general differ by a factor of $1.5-3$. And the timing of the seasonal cycle is estimated within 30 days for Asia, Australia and the maritime continent. For North America and Europe the estimates differ by 60 – 80 days. Note that the outliers for the timing are not necessarily the reanalyses. For North America it is GAUGERT and for Europe it is CMORPH that place the maximum of the annual cycle much earlier in the year than the other estimates. South America and Africa have two maxima in the seasonal cycle, and there is disagreement among data sets on which maximum dominates.

The temporal evolution of global land-averaged precipitation rates on annual, monthly and weekly timescales are shown in Fig. **??**. The interannual variability that can be seen in the annual means is somewhat consistent among most data sets, although there appears to be an offset of 0.5 – 1mm d^{-1} between the estimates (Fig. **??**a), this decreases to 0.3mm d^{-1} when anomalies from the seasonal cycle are considered (not shown). The outliers for annual averages are PERSIANN and to a lesser degree CMORPH. CFSR appears to have a positive trend from 2001 to 2010 not seen in the other estimates; this trend is mostly due to trends over South America and Africa (not shown) and can be related to the changing observing system (**?**). Previous studies have shown that precipitation from reanalyses that assimilate moisture from satellite observations are strongly affected by changes in the observing system and result in spurious trends in the precipitation estimates (**?**). PERSIANN has anomalously high rain rates from late 2006 to early 2007 and anomalously low rate in late 2005 and early 2008 (Fig. **??**b). Over the global ocean the differences among annual averages are larger, up to 2mm d^{-1}, and the reanalyses have a small but significant upward trend not seen in the GPCP, PERSICDR and TRMM estimates (not shown). PERSIANN in contrast has a negative trend over the ocean.

The timing of the seasonal cycle over North America is captured more or less consistently by all estimates (Fig. **??**b), but the amplitude is not. CMORPH and PERSIANN underestimate winter precipitation rates relative to other analyses by up to 1mm d^{-1} on monthly time scales, while ERAI under-estimates summer precipitation rates. On weekly time scales the differences can be as large as 3mm d^{-1} in the winter, with CMORPH and PERSIANN estimating < 0.5mm d^{-1} and all other estimates averaging between $2.5 - 3$mm d^{-1} (Fig. **??**c). This is a very large range for an area of this size and a weekly average. This is a known issue with CMORPH and PERSIANN. Several studies have shown that wintertime precipitation is severely underestimated in these products for different regions in the northern midlatitudes (**???**). Relative differences over North America in the summer are of the same order as over the maritime continent, even though total amounts are much larger over the maritime continent.

Correlations of the time series of continental mean precipitation anomalies reveal large positive correlations on annual, monthly and daily time scales for some data sets, TRMM, PERSICDR and GPCP in particular (Table **??**). For other data sets the correlations are not significantly different from zero, even for annual averages (GPCP and PERSIANN or CMORPH), indicating potential long-term differences in the continental scale water budgets associated with the different data sets that would need to be balanced by evaporation or runoff. Results for reanalyses are mixed. Correlations on annual timescales are < 0.3 for all reanalyses over North America, but > 0.9 over Europe, the maritime contitent and Australia. Meanwhile, correlations are fairly high for both monthly and daily timescales.

The low correlations of large scale (continental to global) annual averages of precipitation estimates indicate that the estimates differ in their interannual variability. Imbalances on these scales in estimates of an important component of the global water cycle affect our ability to close the budget (**??**). Global land differences on annual time scales are about 0.8mm d^{-1} for the observational estimates. This translates to differences of up to 23.2W m^{-2}, which is very large compared to the global land latent heat flux of 38.5W m^{-2} estimated by **?**. Including the reanalyses increases the offset to 1mm d^{-1}.

d. Distributions

Fig. **??** shows the area-averaged seasonal distributions for North America. The general behavior of these distributions is very similar for the other continental areas. The log-log plot shows curves with two distinct slopes, positive for low rain rates and negative for higher rain rates. The transition between these slopes is more abrupt in the summer and more gradual in the winter months for North America. For Africa and the maritime continent, the transition is abrupt for all months (not shown). This relationship appears to hold for all continental areas during the summer months when precipitation tends to be in a more convective regime, which leads us to speculate that the manner of transition between slopes could be related to the dominant precipitation regime (large-scale vs. convective). While the location of where the slopes in the log-log plot change is around 0.5mm h^{-1} for all seasons and regions, the slopes are very variable between months, data sets and regions.

At the lowest rain rates, JRA55, MERRA and CMORPH, have a positive bias, with lower rain rates being more common than in other reanalyses or the precipitation data sets. This is consistent with all other continental areas except the maritime continent, where GAUGERT and CMORPH have a positive bias at low rain rates and ERAI and MERRA have a low bias. The distributions over the maritime continent have the largest spread among the data sets. The bulk of the distribution is between $0.01 - 1$mm h^{-1}, with the peak in the distribution shifting between 0.015mm h^{-1} in the winter and 0.5mm h^{-1} in the summer for North America (Fig. **??**c). In general, reanalyses, and MERRA in particular, dominate the distribution at these rates. For midlatitude continental regions, CMORPH, and PERSIANN to a lesser degree, are a lot less likely than other products to have precipitation occur at the intermediate rates $0.01 - 1$mm h^{-1}. This is likely related to the fact that these are satellite-only products that have issues with detecting precipitation over snow-covered ground. Fig. **??** examines the differences in the tails of the precipitation distributions. Overall reanalyses tend to not produce very high rain rates. This could be because of the grid area vs. point estimate, the convective parameterizations used, or the relatively large grid size. For North America in the winter TRMM has the highest rain rates and highest probability of high rates occurring (Fig. **??**a). In the summer (Fig. **??**c) the satellite only estimates dominate at the highest rain rates. For other regions ERAI dominates the tails in the winter in South America and all year in Africa (not shown).

The satellite-only products, CMORPH and PERSIANN, tend to accentuate the tail of the distribution during summertime convective precipitation regimes. During months when precipitation is dominated by synoptic systems or when the ground is covered in snow (e.g. Europe in the winter months) the tails of the distributions of CMORPH and PERSIANN are even lower than the reanalyses.

A different way to compare the data sets is through the distribution of the rain amount by rain rate (Fig. **??**). Precipitation amount distributions tend to be skewed in a logarithmic plot, with a long tail towards lower rain rates. Rain rates below 0.01mm h^{-1} are very common, but the actual rain amount from precipitation at these rates does not add up to much. During the winter months (Fig. **??**a), the distributions for CMORPH and PERSIANN are much flatter, and the mean total precipitation amount of CMORPH in DJF is 14mm, whereas it is 55mm for GPCP and 68mm for CFSR. That is a difference of almost 500% for the mean monthly total estimate. Excluding CFSR which has been shown to overestimate moisture transport from ocean to land and where at least some of the precipitation over land is due to the analysis increment (**?**), there is still a factor of 4 difference. On the other hand, in summer (Fig. **??**c), CMORPH and PERSIANN have many high rain rate events compared to the other estimates, and their monthly mean totals are correspondingly higher than the other estimates. One thing to note about the reanalysis estimates is that the rain amount distributions tend to be narrower than the satellite and rain gauge estimates. This is most obvious for ERAI (Fig. **??**c) and becomes more severe for reanalyses with a coarser spatial resolution (not shown), highlighting the fact that reanalyses only resolve a narrow band of rain rates.

5. Summary and Discussion

A comparison of several global precipitation estimates and reanalyses was performed on a range of temporal and spatial scales. Only data sets with daily or higher temporal resolution were considered. We found that while patterns of means and variance were largely consistent among data sets, the differences in means and variances between the data sets were often as large as the analyzed means and variances themselves.

Correlations among the precipitation estimates averaged over continental areas varied significantly. GPCP, TRMM and PERSICDR were very highly correlated. This was by construction on monthly and annual time scales, since all three data sets are bias corrected to monthly satellite - rain gauge analyses, but also held for daily averages. Correlations of the satellite-only products, PERSIANN and CMORPH, with GPCP were not significantly different from zero even for annual averages. Reanalyses had high correlations with GPCP on monthly time scales, but the results were mixed for annual averages. Correlations between reanalyses and GPCP were found to be larger than 0.9 over Europe, the maritime continent and Australia, but less than 0.3 over North America. This is noteworthy, because North America is one of the best observed regions in the world where the potential for constraining reanalyses with observations is high.

Distributions of precipitation rates and amounts showed that satellite-only estimates PERSIANN and CMORPH underestimated wintertime precipitation in midlatitudes, while overestimating summertime precipitation in midlatitudes. Reanalyses tended to precipitate over too narrow of a range of rain rates when compared to observational estimates, although some of the reanalyses (JRA55 and MERRA) estimate mean monthly totals in the same range as PERSIANN and CMORPH in the summer. Reanalyses tended to precipitate over too narrow of a range of rain rates when compared to observational estimates, although mean monthly totals of JRA55 and MERRA in the summer were in the same range as PERSIANN and CMORPH. The difference (at least for North America) is that the bulk of the rain in the satellite-only estimates PERSIANN and CMORPH comes from high rain rates $> 2\text{mm h}^{-1}$, while JRA55 overestimation occurred at rain rates around 0.8mm h^{-1} and for MERRA at around 0.4mm h^{-1}.

Average spread among data sets was computed for each grid point, and is defined as the average of the standard deviation of anomalies from the seasonal cycle. Spread among data sets differed between reanalyses and satellite estimates (Fig. **??**). Spread among reanalyses was found to be larger in the tropics and smaller in midlatitudes when compared to the spread among satellite estimates. This is likely related to precipitation in midlatitudes being driven mainly by the large-scale flow, while convective precipitation dominates in the tropics. Reanalyses do well in representing mid-latitude large-scale circulation patterns and this results in higher consistency across reanalyses in the mid-latitudes. In the tropics convective parameterizations were likely responsible for the bulk of the precipitation in reanalyses; these parameterizations differed widely among reanalyses and so did the results.

Systematic differences were found in the global precipitation estimates considered in this study. Users of these estimates need to be aware of these biases and their use as a ground truth should be limited to regimes, seasons, or regions the products have been shown to perform well for. For example, CMORPH and PERSIANN, designed to represent the instantaneous variability in precipitation, performed well in the tropics, but overestimated summertime convective precipitation and underestimated wintertime precipitation in midlatitudes. This suggests that the performance of CMORPH or PERSIANN in midlatitude regions always needs to be assessed for the region and season of interest prior to using these estimates. Reanalyses reflect the systematic errors of the global circulation models used to provide the forecast background. There is a clear bias of the reanalyses' annual and monthly means compared to the observational estimates. However, while we showed here that large scale (continental to global) annual averages of precipitation estimates differ in their interannual variability, variability estimated by reanalyses on monthly timescales tends to be consistent with the observational estimates (as seen from the high correlations). This suggests that studies interested mainly in the variability of precipitation may have a more reliable foundation in using reanalyses than studies investigating the energy and water budgets.

In summary, any study using precipitation estimates based on observations or reanalyses should take into account the uncertainty associated with the precipitation estimate. There is no one global precipitation product that is better than all the others for all applications. The most suitable product changes with intended application, location and season. Therefore, care needs to be taken when choosing a product for a specific application, to ensure that the product has the capability to yield useful results. Given the uncertainty inherent in any precipitation estimate it is an asset to have several products based on different approaches available to compare and estimate that uncertainty. In some ways precipitation estimates from satellite and reanalyses have the opposite problem. Satellite estimates perform well in regions and seasons with convective precipitation, while reanalyses are better at large scale precipitation in the northern midlatitudes. Precipitation estimates that incorporate both satellite and ground-based measurements such as GPCP, and indirectly TRMM and PERSICDR, tend to lie in between the other estimates both in terms of the distributions and the average rain rates. Incorporating ground radar in precipitation estimates where available can be expected to have a positive impact on the accuracy of the estimates. Including data from diverse sources (multiple satellites and retrieval channels, rain gauge, radar) appears to help with reducing errors and enhances reliability. Extending the rain gauge network to data sparse regions, in particular over oceans, will likely have a large impact on constraining at least global mean precipitation estimates. Unfortunately, this is impractical and costly. A more practical approach may be to combine precipitation estimates from several different data sources based on their respective strengths.

Acknowledgments. GPCP daily data are available courtesy of NASA at http://precip.gsfc.nasa.gov. TRMM is available at http://disc.gsfc.nasa.gov/precipitation. The PERSIANN precipitation product is available at http://chrs.web.uci.edu/persiann/data.html. NOAA CPC Morphing Technique (CMORPH) Global Precipitation Analyses is available at the Research Data Archive at the National Center for Atmospheric Research, Computational and Information Systems Laboratory http://rda.ucar.edu/datasets/ds502.0/. The PERSIANN CDR (PERSICDR) used in this study was acquired from NOAA's National Centers for Environmental Information (http://www.ncdc.noaa.gov). This CDR was originally developed by Soroosh Sorooshian and colleagues for NOAA's CDR Program. GAUGE data was obtained from the Climate Prediction Center at http://ftp.cpc.ncep.noaa.gov/precip/CPC_UNI_PRCP/GAUGE_GLB/. ERA-Interim data provided courtesy ECMWF and the Research Data Archive at the National Center for Atmospheric Research. The CFSR dataset used for this study is provided from the Climate Forecast System Reanalysis (CFSR) project carried out by the Environmental Modeling Center (EMC), National Centers for Environmental Prediction (NCEP). The (JRA55) dataset used for this study is provided from the Japanese 55-year Reanalysis project carried out by the Japan Meteorological Agency (JMA). MERRA was developed by the Global Modeling and Assimilation Office and supported by the NASA Modeling, Analysis and Prediction Program. Source data files can be acquired from the Goddard Earth Science Data Information Services Center (GES DISC).

LIST OF TABLES

TABLE 1. List of precipitation estimate data sets. Sources are geostationary infrared (Geo-IR), microwave (MW), gauges, or reanalyses.

Name	Source	Temporal resolution	Spatial resolution	Reference
TRMM	Geo-IR; MW from SSM/I,TMI,	1998 - 2012,	49°S - 49°N	**?**
	AMSU, AMSR; gauges	3 hourly	0.25°	**?**
CMORPH	Geo-IR; MW from SSM/I,TMI,	2003 - 2013,	59°S - 59°N	**?**
	AMSU, AMSR;	3 hourly	0.25°	**?**
PERSIANN	Geo-IR; MW from TMI	2001 - 2013,	59°S - 59°N	**??**
		3hourly	0.25°	**?**
PERSICDR	Geo-IR; MW from TMI (for training)	1983 - 2013,	60°S - 60°N	**??**
	SSM/I; IR; gauges	daily	0.25°	**?**
GPCP	Geo-IR; AVHRR low-earth-orbit IR,	1997 - 2013,	global, 1°	**?**
	SSM/I; gauges;	daily		
	TOVS (poleward of 40S-40N)			
GAUGE	gauges	1979 - 2005, daily	global land, 0.5°	**??**
GAUGERT	gauges	2006 - 2013, daily	global land, 0.5°	**??**
JRA55	Reanalysis	1979 - 2013,	global, gaussian 0.5625°	**?**
MERRA	Reanalysis	1979 - 2013	global, 0.5° x $2/3$°	**?**
CFSR	Reanalysis	1979 - 2010	global, 0.5°	**?**
ERAI	Reanalysis	1979 - 2013	global, 0.75°	**?**

TABLE 2. Description of the metrics used in the analysis. $P(x,y,d,m,yr)$ is precipitation at longitude x, latitude y, day d, month m, and year yr. N_m is the total number of days in month m, $m = 1,...,12$. N_A is the number of grid points in region A with $(x_i, y_j) \in A$. w_j are the weights that account for changing area of the grid box with latitude. $P_1,...,P_{N_d}$ are the different data sets, with N_d the total number of data sets. M is the mean of all the precipitation data sets.

Metric	
Monthly mean	$\bar{P}(x,y,m) = \frac{1}{N_m}\sum_{yr=1}^{N}\sum_{k=1}^{N_{my}} P(x,y,d_k,m,yr)$
Monthly variance	$\sigma^2(x,y,m) = \frac{1}{N_m}\sum_{yr=1}^{N}\sum_{k=1}^{N_{my}} (P(x,y,d_k,m,yr) - \bar{P}(x,y,m))^2$
Difference	$D(x,y,m) = \bar{P}(x,y,m) - \bar{Q}(x,y,m)$
Percentage difference	$D(x,y,m) = \frac{\bar{P}(x,y,m)-\bar{Q}(x,y,m)}{\bar{P}(x,y,m)} * 100$
Spatial average	$P_A(d,m,yr) = \frac{1}{N_A}\sum_{i=1}^{N_{xA}}\sum_{j=1}^{N_{yA}} w_j P(x_i,y_j,d,m,yr)$
Spread among data sets	$\sigma_P(x,y) = \frac{1}{N_t}\sum_{k=1}^{N_t}\sqrt{\frac{1}{N_d}\sum_{d=1}^{N_d}(P_d(x,y,t_k) - M(x,y,t_k))^2}$

TABLE 3. Description of continental regions used in the analysis. Only points over land inside the domains are used. Also shown are the amplitude (mm d^{-1}) of the area averaged mean annual cycle for 2006-2012 and the phase (the day of the year the maximum occurs). These are given for all data sets in the order (TRMM, GPCP, CMORPH, PERSIANN, PERSICDR, GAUGERT, JRA55, MERRA, CFSR, ERAI). The minimum and maximum are highlighted in bold.

Region	lon-lat	Amplitude	Phase
North America	165°W - 50°W	(1.49, **1.18**, 1.25, 1.2, **1.18**,	(270, **274**, 257, 251, **274**,
	15°N - 49°N	1.42, 1.5, **1.54**, 1.37, 1.2)	**188**, 265, 256, 267, 270)
South America	90°W - 30°W	(1.35, 1.29, 1.43, 1.98, 1.3,	(**74**, 71, 315, **306**, 69,
	49°S - 15°N	**3.45**, 1.32, 1.2, 1.7, **1.12**)	54, 48, 47, 327, 338)
Europe	15°W - 50°E	(1.57, 1.52, 0.66, **0.63**, 1.49,	(311, 329, **285**, 301, 332,
	30°N - 49°N	0.84, 1.3, 0.91, **1.65**, 1.08)	314, 316, 328, **347**, 322)
Africa	20°W - 50°E	(0.67, **0.55**, 0.78, 0.91, 0.58,	(98, **85**, 108, 101, 87,
	35°S - 30°N	**0.94**, 0.77, 0.87, 0.71, 0.78)	**227**, 98, 103, 240, 95)
Asia	50°E - 150°E	(4.03, 3.75, 3.54, 3.31, 3.83,	(200, 200, **188**, 195, 199,
	5°N - 49°N	**3.1**, **5.1**, 4.48, 4.44, 3.39)	203, 201, **209**, 201, 205)
Maritime Continent	90°E - 165°E	(3.22, **2.98**, 3.37, **4.44**, **2.98**,	(363, 1, 363, 365, 1,
	10°S - 5°N	4.23, 4.21, 3.09, 3.53, 3.19)	**351**, 363, 2, **14**, 366)
Australia	110°E - 155°E	(3.27, 2.89, 3.3, **3.96**, 2.95,	(29, 36, 22, **21**, 35,
	49°S - 10°S	3.21, 3.69, 3.09, 2.49, **2.18**)	**37**, 28, 28, **21**, **37**)

TABLE 4. Correlations between GPCP and all other data sets for annual, monthly and daily mean time series. Correlations are computed for common time period 2003-2010 with the annual cycle removed. Correlations significant at the 90% level are bold.

	TRMM	CMORPH	PERSIANN	PERSICDR	JRA55	MERRA	CFSR	ERAI
				Annual				
North America	**0.84**	0.67	0.05	**0.97**	0.13	0.28	-0.10	0.24
South America	**0.99**	-0.16	-0.27	**1.00**	**0.83**	0.62	0.48	**0.81**
Europe	**0.96**	0.02	-0.29	**0.99**	**0.94**	**0.93**	**0.91**	**0.92**
Africa	**0.98**	-0.05	**0.71**	**1.00**	0.62	**0.81**	0.37	0.59
Asia	**0.99**	0.04	0.32	**1.00**	**0.77**	**0.77**	0.47	0.64
maritime continent	**0.99**	0.7	0.66	**1.00**	**0.94**	**0.91**	**0.96**	**0.95**
Australia	**0.99**	**0.94**	0.32	**1.00**	**0.94**	**0.90**	**0.99**	**0.90**
				Monthly				
North America	**0.98**	**0.56**	**0.38**	**0.98**	**0.88**	**0.87**	**0.84**	**0.84**
South America	**0.99**	**0.27**	0.14	**0.98**	**0.80**	**0.66**	**0.54**	**0.70**
Europe	**0.96**	**0.39**	0.16	**0.99**	**0.95**	**0.95**	**0.93**	**0.95**
Africa	**0.98**	**0.24**	**0.47**	**1.00**	**0.63**	**0.67**	**0.52**	**0.65**
Asia	**0.98**	**0.26**	**0.28**	**0.99**	**0.84**	**0.81**	**0.70**	**0.82**
maritime continent	**0.99**	**0.86**	**0.76**	**1.00**	**0.96**	**0.96**	**0.96**	**0.96**
Australia	**0.99**	**0.88**	**0.57**	**1.00**	**0.88**	**0.91**	**0.93**	**0.85**
				Daily				
North America	**0.78**	**0.68**	0.01	**0.91**	**0.71**	**0.57**	**0.68**	**0.66**
South America	**0.86**	**0.78**	-0.00	**0.90**	**0.72**	**0.63**	**0.63**	**0.65**
Europe	**0.80**	**0.55**	0.01	**0.89**	**0.67**	**0.62**	**0.66**	**0.64**
Africa	**0.88**	**0.78**	**-0.05**	**0.96**	**0.73**	**0.63**	**0.54**	**0.65**
Asia	**0.86**	**0.79**	**-0.06**	**0.96**	**0.81**	**0.66**	**0.77**	**0.75**
maritime continent	**0.89**	**0.86**	0.01	**0.97**	**0.84**	**0.77**	**0.83**	**0.82**
Australia	**0.92**	**0.87**	-0.03	**0.99**	**0.80**	**0.76**	**0.80**	**0.76**

LIST OF FIGURES

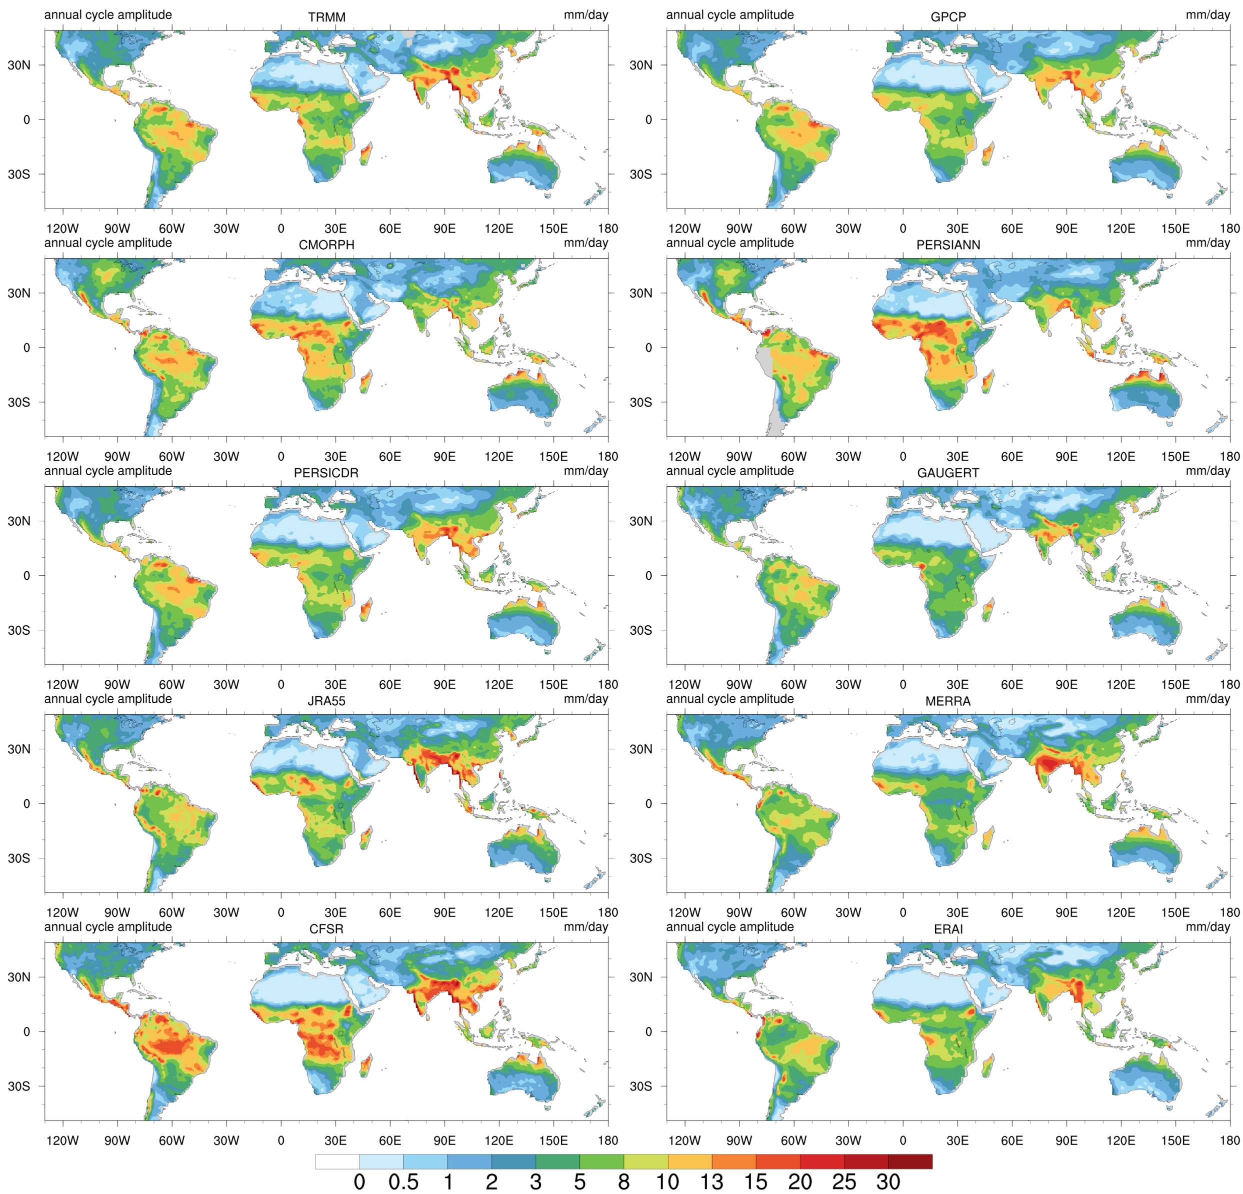

FIG. 1. Annual cycle amplitude in mm d^{-1}. The annual cycle is computed as the first 4 harmonics of the mean daily seasonal cycle from 2006 – 2012. The amplitude is the difference between the minimum and maximum of the annual cycle.

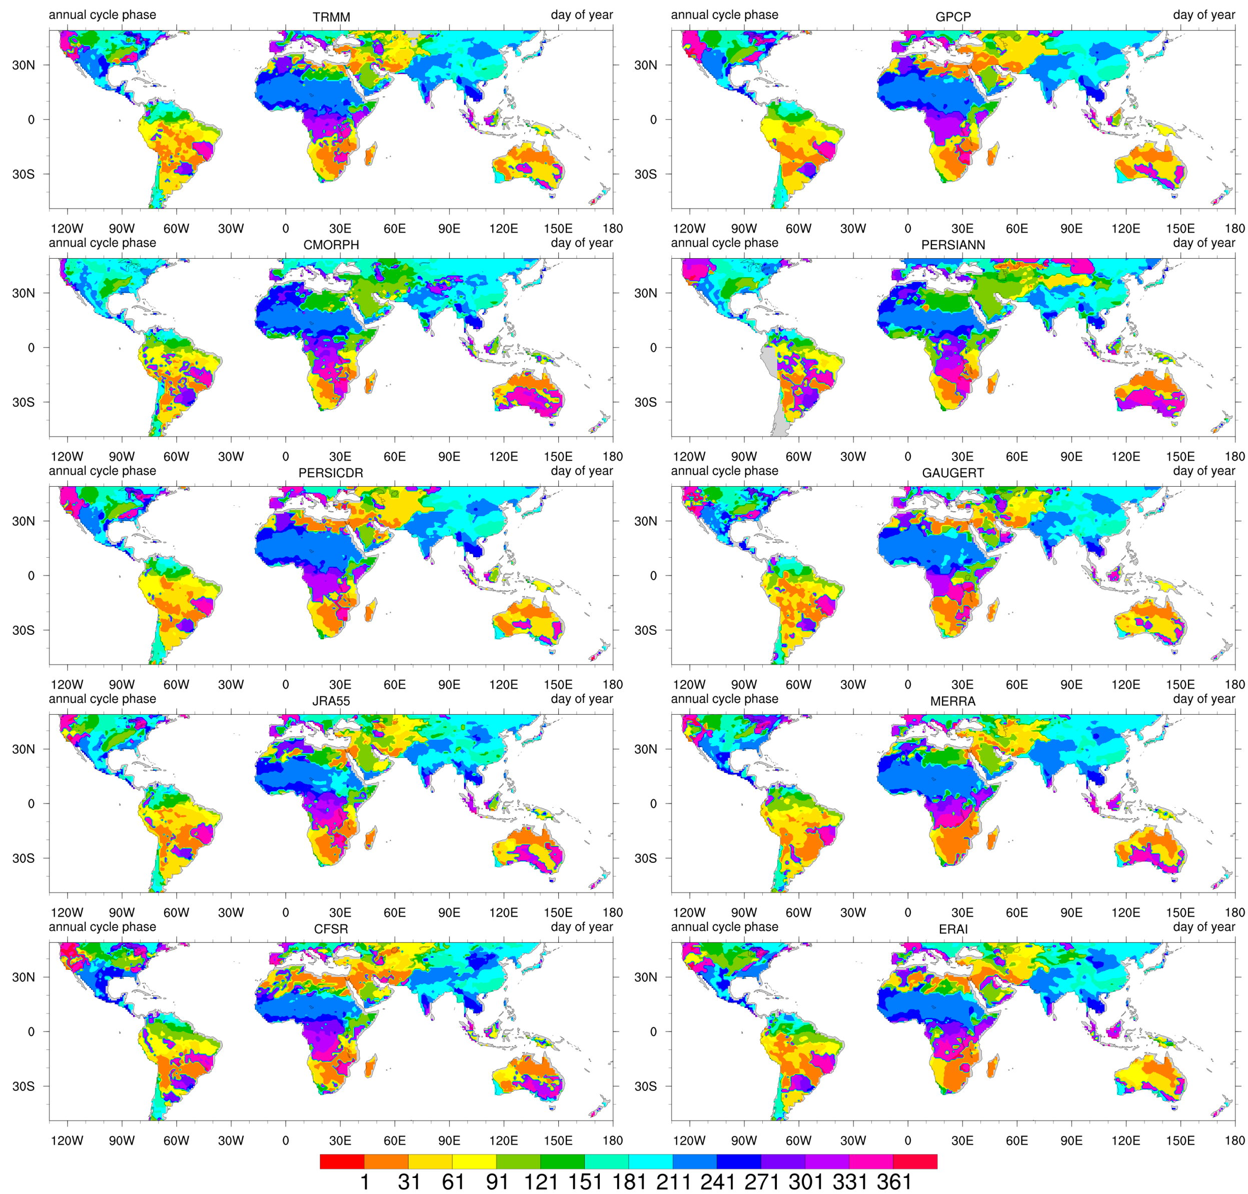

FIG. 2. Annual cycle phase in day of year. The annual cycle is computed as the first 4 harmonics of the mean daily seasonal cycle from 2006 – 2012. The phase is the day of the year the maximum of the annual cycle is achieved.

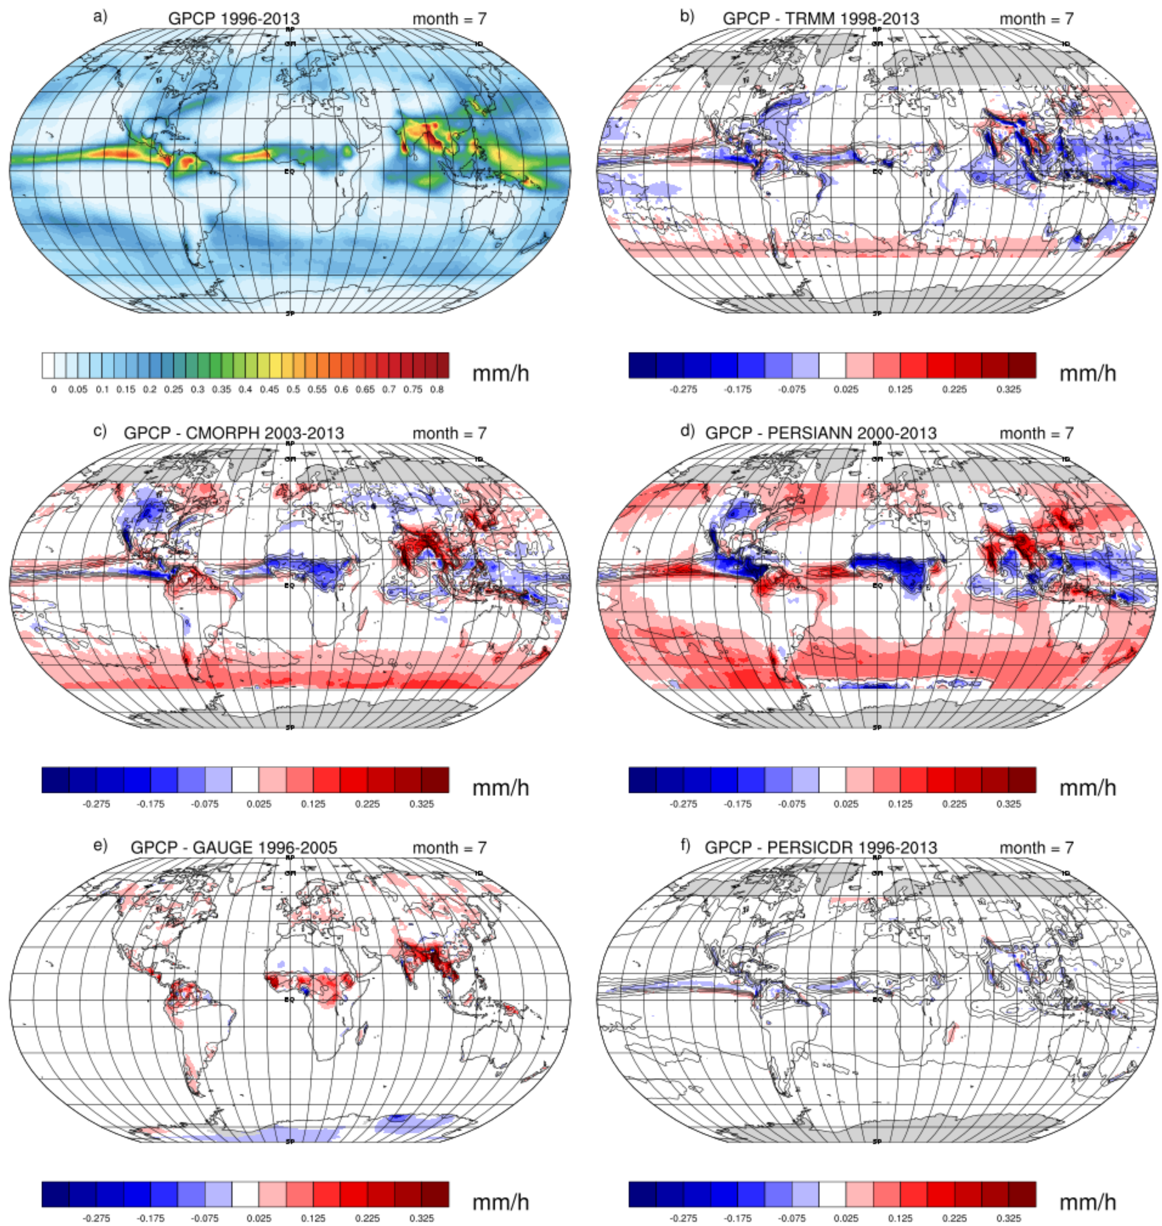

FIG. 3. Monthly long term means of precipitation for July. a) mean for GPCP. b)-f) the difference between GPCP mean and the respective data set mean for the period is indicated in shading, contours show the mean monthly precipitation for the respective data set. Contour levels go from 0 to 0.4 by 0.1mm h^{-1}.

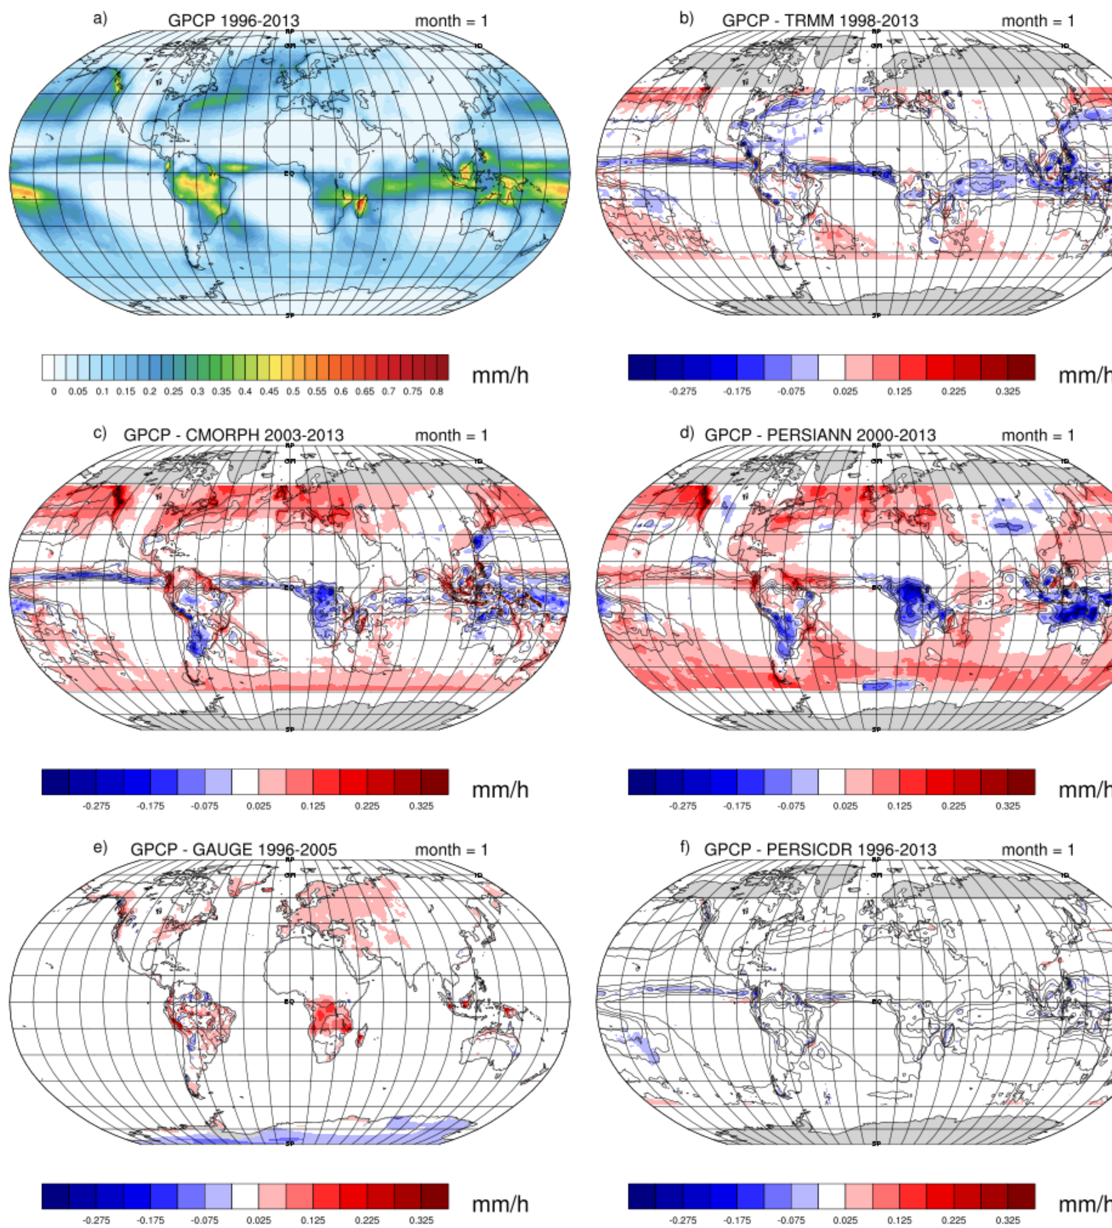

FIG. 4. Same as in Fig. **??**, but for January.

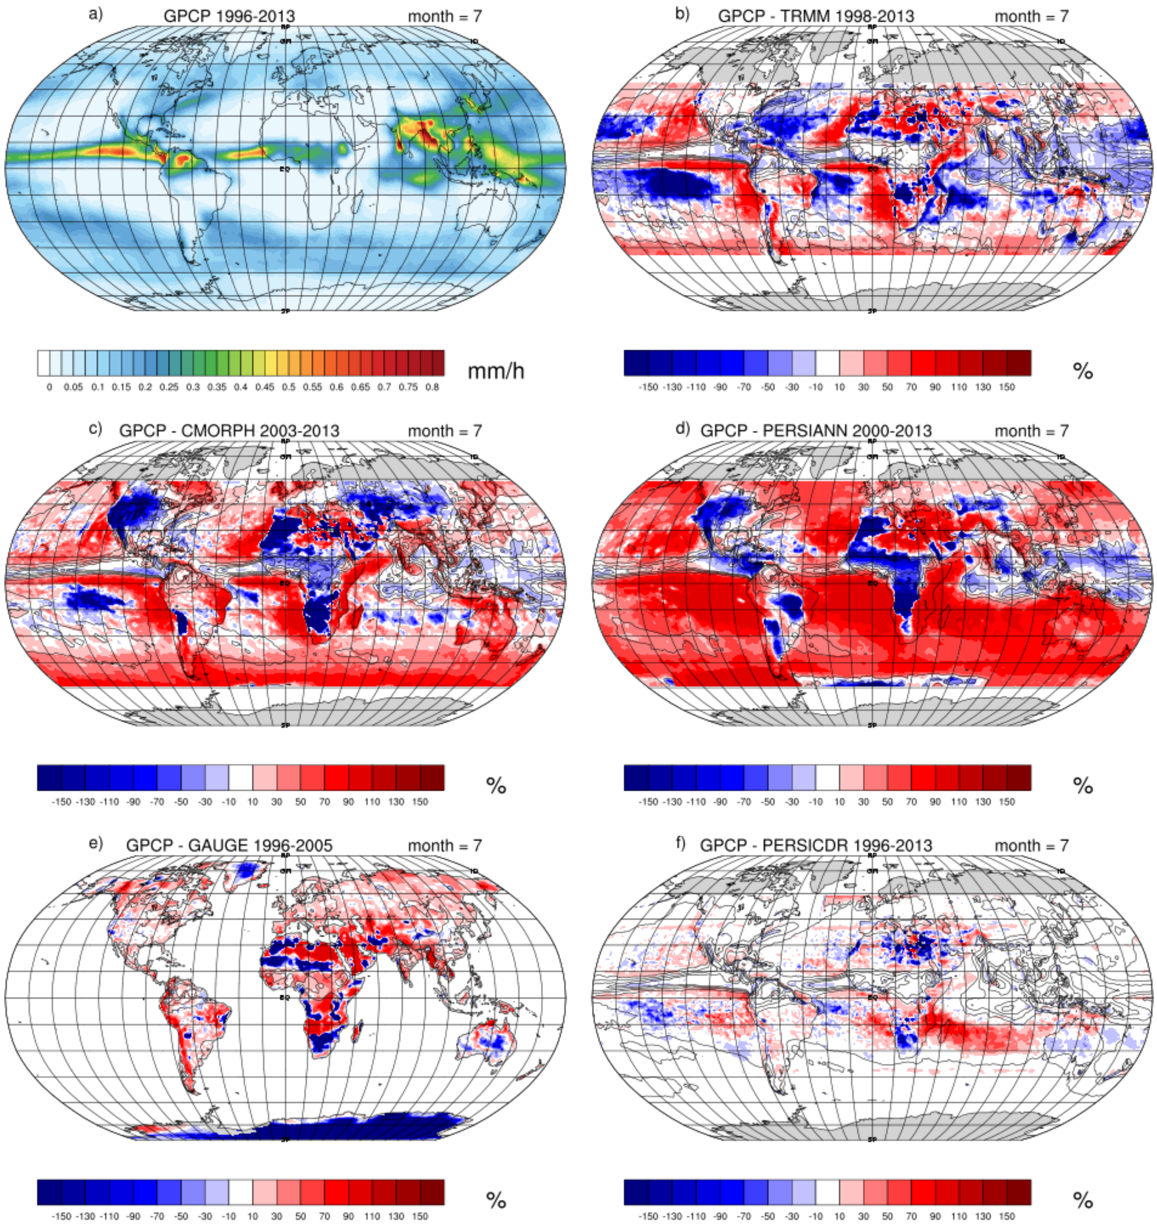

FIG. 5. Monthly long term means of precipitation for July. a) mean for GPCP. b)-f) the percentage difference between GPCP mean and the respective data set mean for the period is indicated in shading, contours show the mean monthly precipitation for the respective data set. Contour levels as in Fig. **??**.

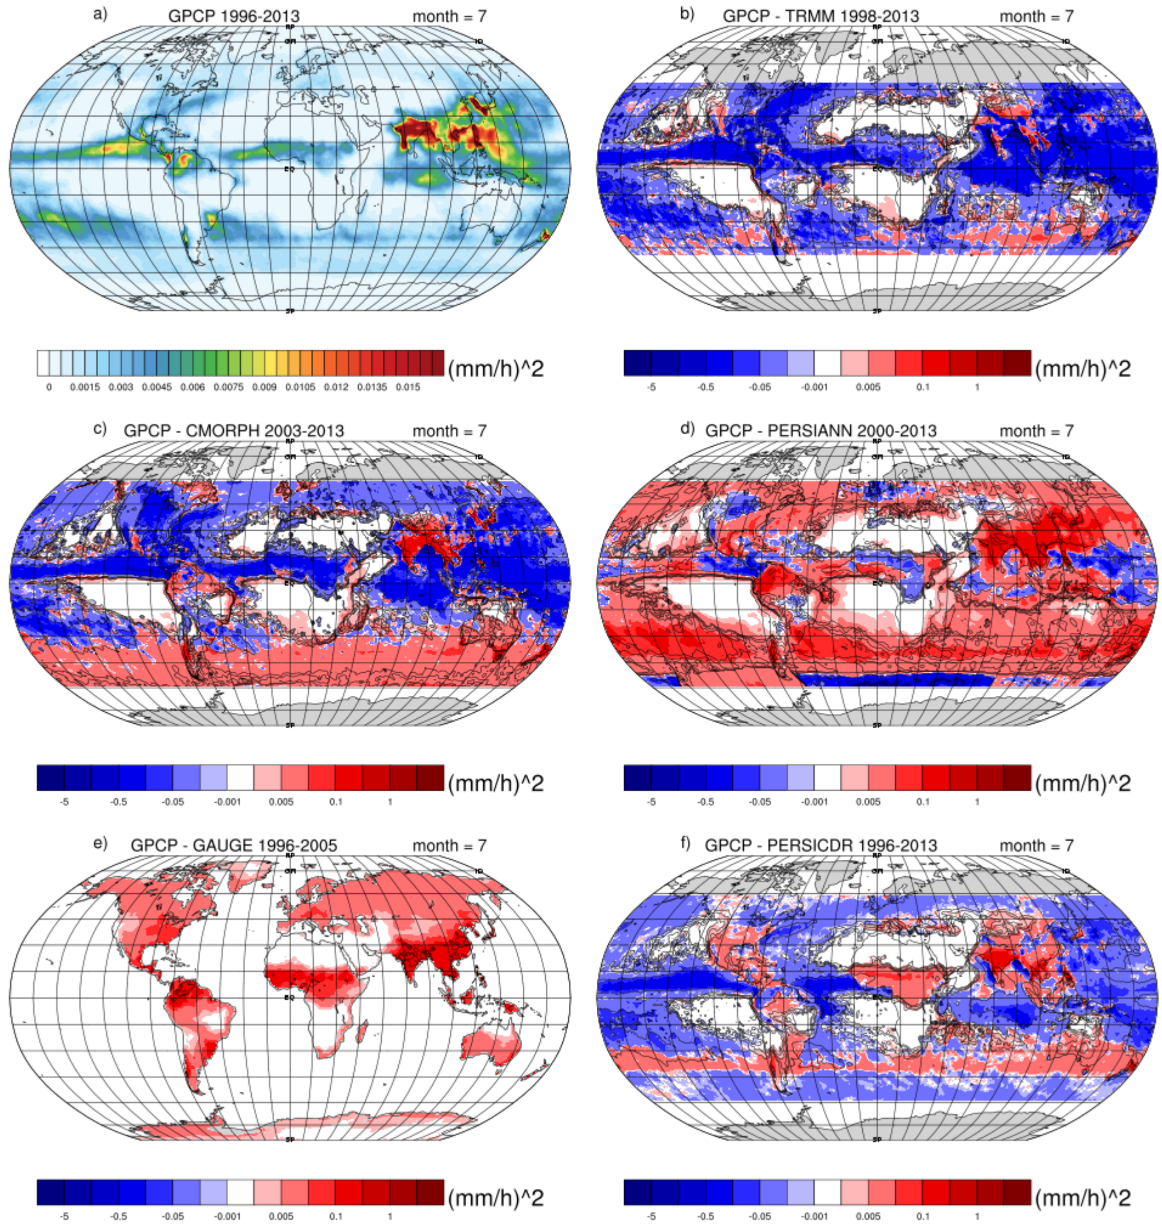

FIG. 6. Monthly mean variance of precipitation for July. a) mean variance for GPCP. b)-f) the difference between the GPCP mean variance and the respective data set mean variance for the period is indicated in shading, contours show the mean monthly precipitation variance for the respective data set. Contour levels are $(0.001, 0.002, 0.005, 0.01, 0.1, 1, 2, 10)$.

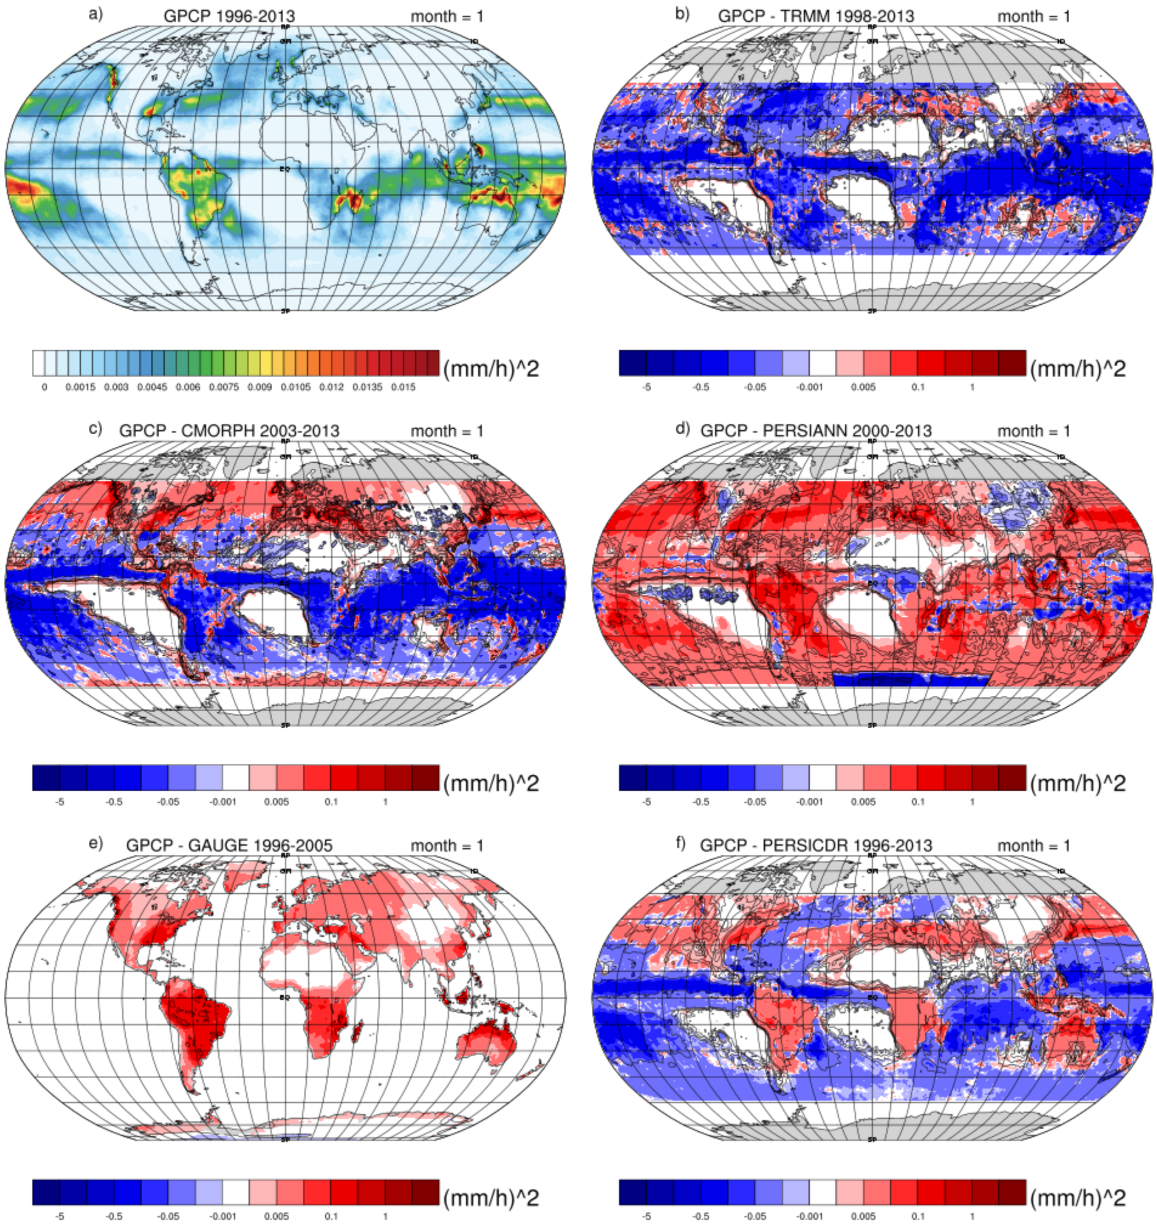

FIG. 7. Same as in Fig. **??**, but for January.

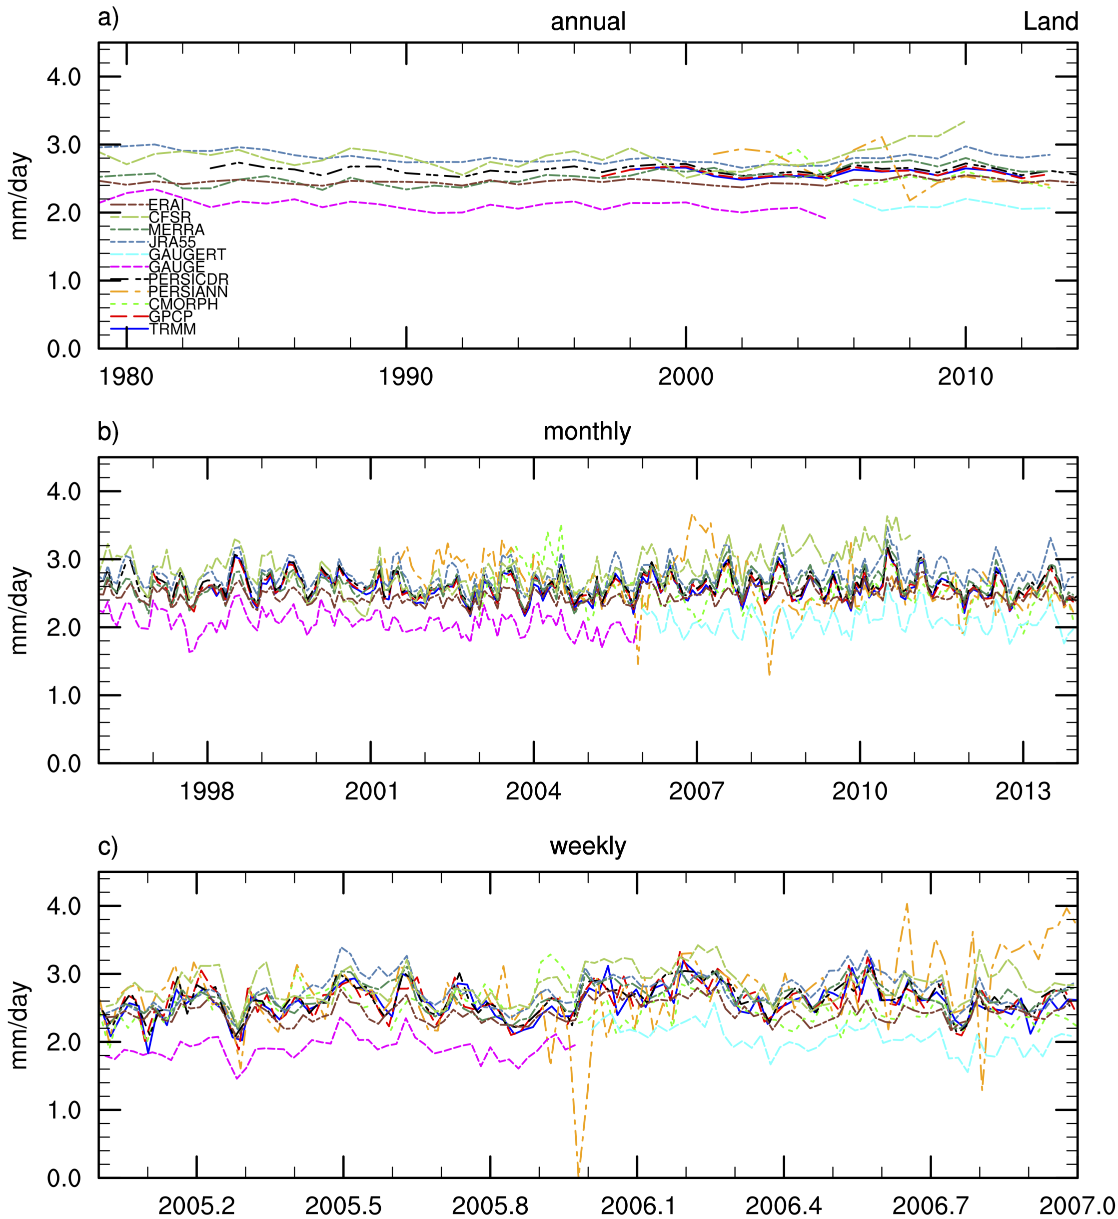

FIG. 8. Time series of rain rates averaged over global land area between 49°N and 49°S for a) annual means, b) monthly means, and c) weekly means.

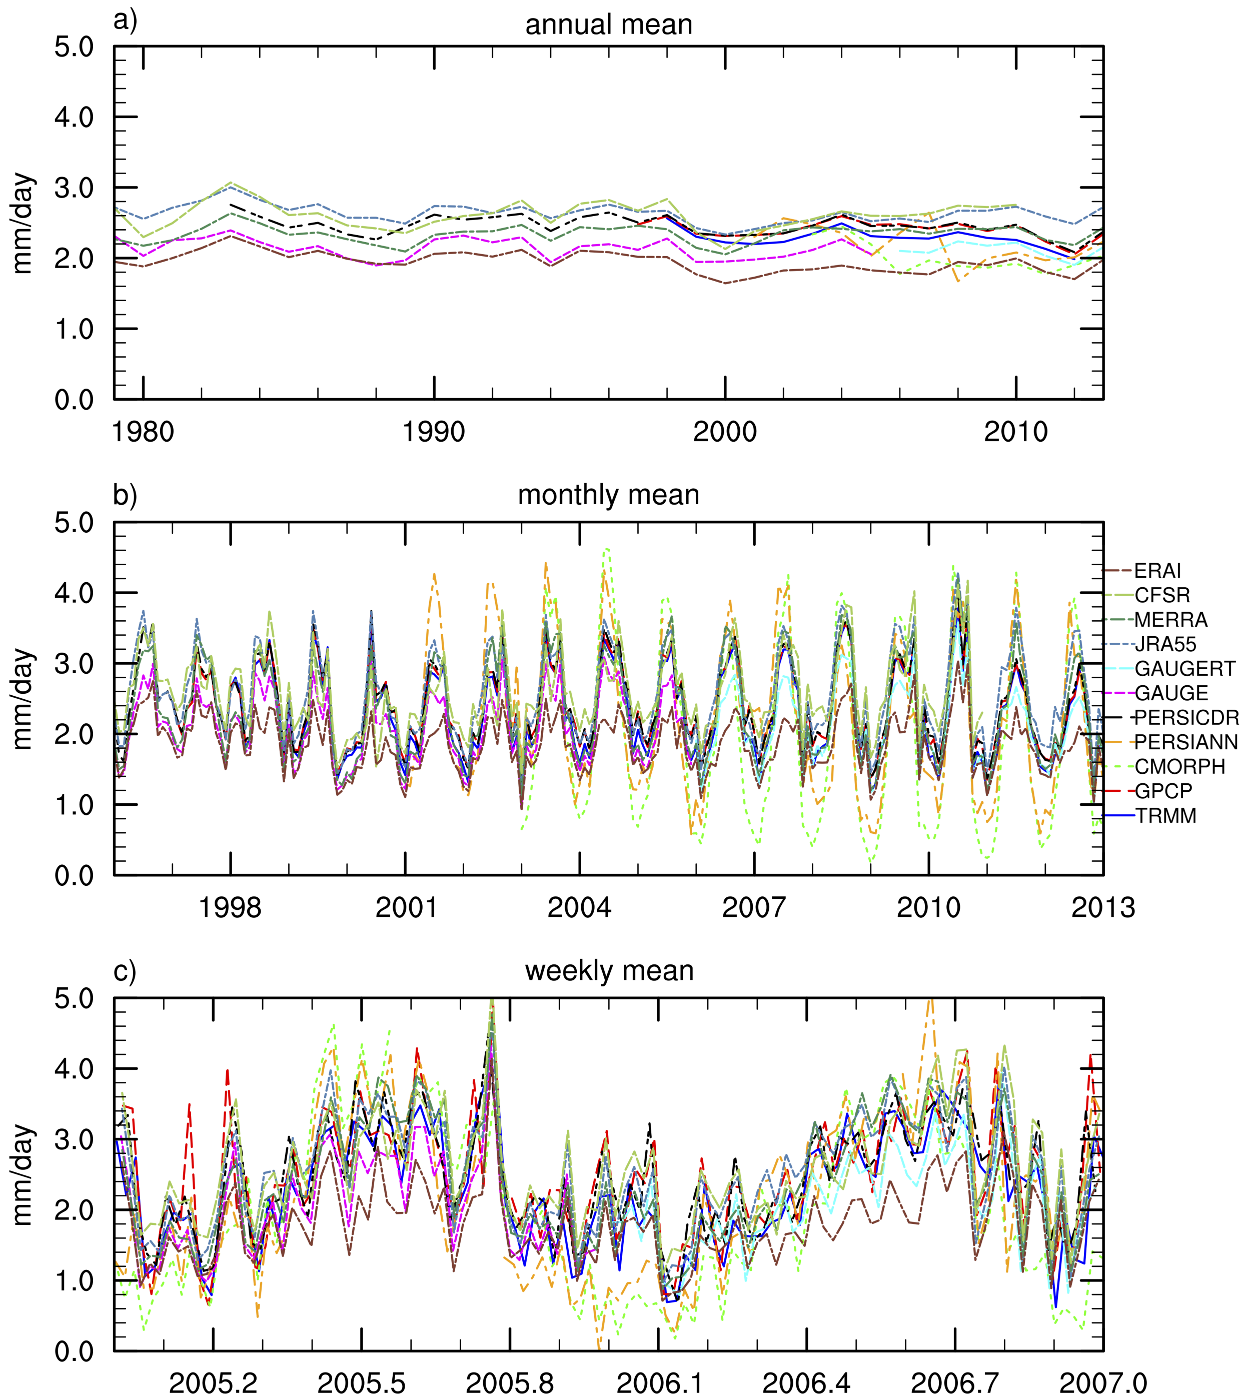

525 FIG. 9. Time series of rain rates averaged over North America land area between 15 – 49°N for a) annual
526 means, b) monthly means, and c) weekly means.

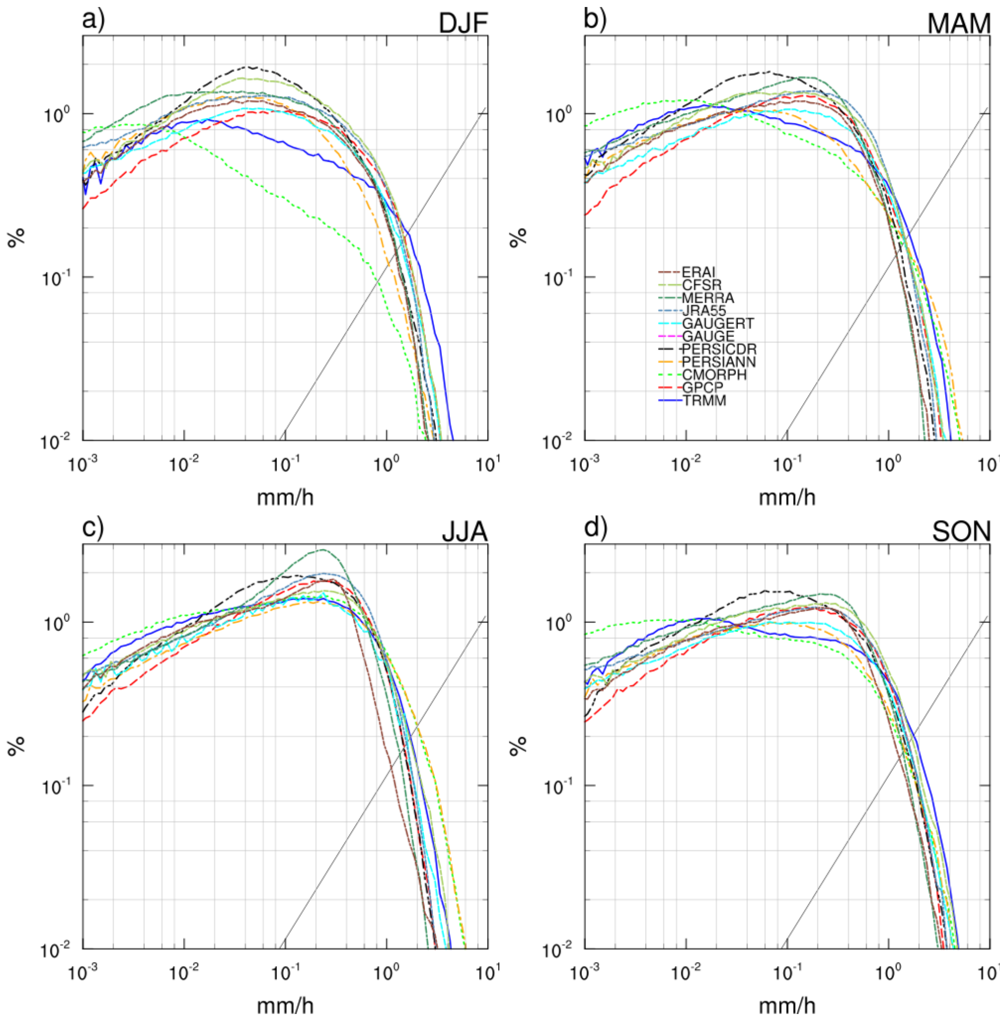

FIG. 10. Distribution of precipitation rate over land area for North America (15°N - 49°N, 195°E - 310°E). Panels a)-d) show the climatological distribution for all seasons for 2006 - 2012. Precipitation rates are binned with logarithmic bin sizes to account for more frequent rain events at low rain rates. The x axis is plotted on a log-scale to compare the bulk of the distribution, not the tails. The black line shows the size of the bin at each precipitation rate. Distributions are computed for each month and grid point separately and then averaged over area and season.

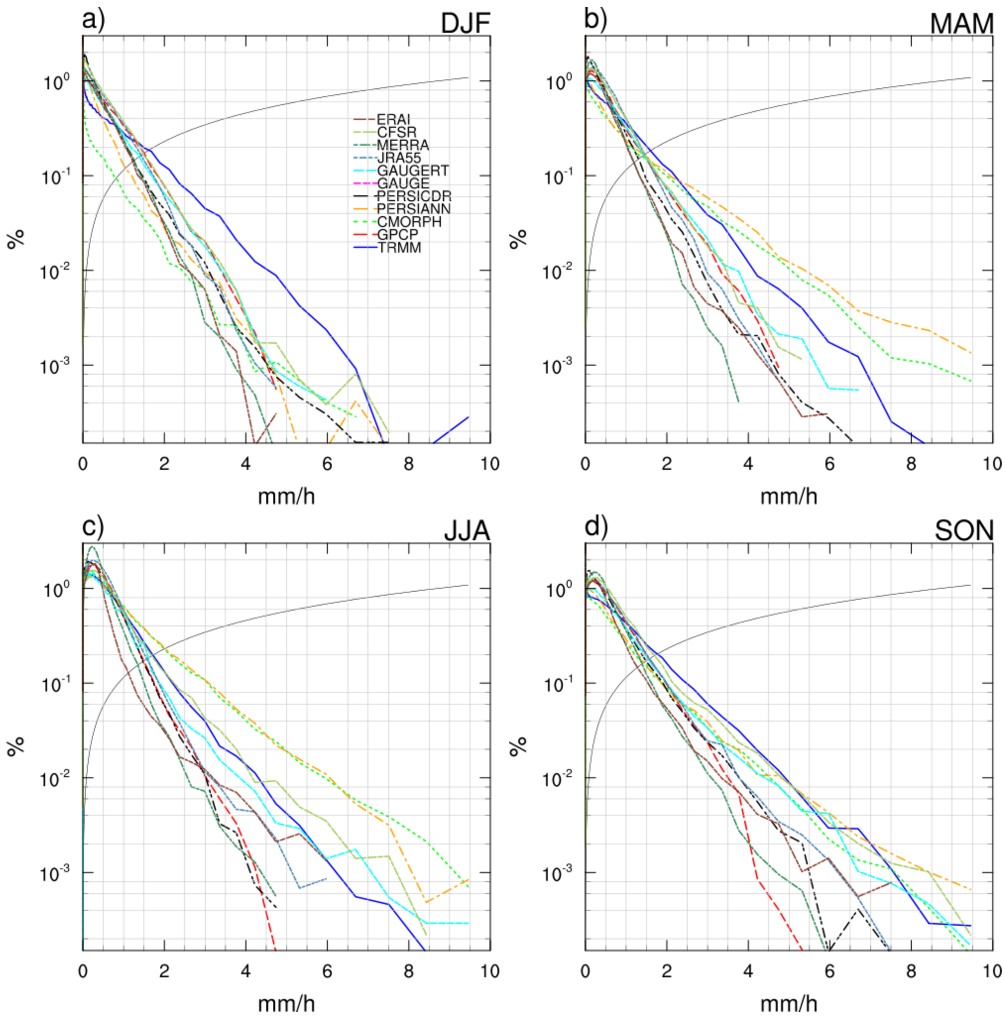

FIG. 11. Distribution of precipitation rate over land area for North America (15°N - 49°N, 195°E - 310°E). As in Fig. **??**, except that the x axis is plotted on a linear scale to facilitate comparison of the tails of the distributions.

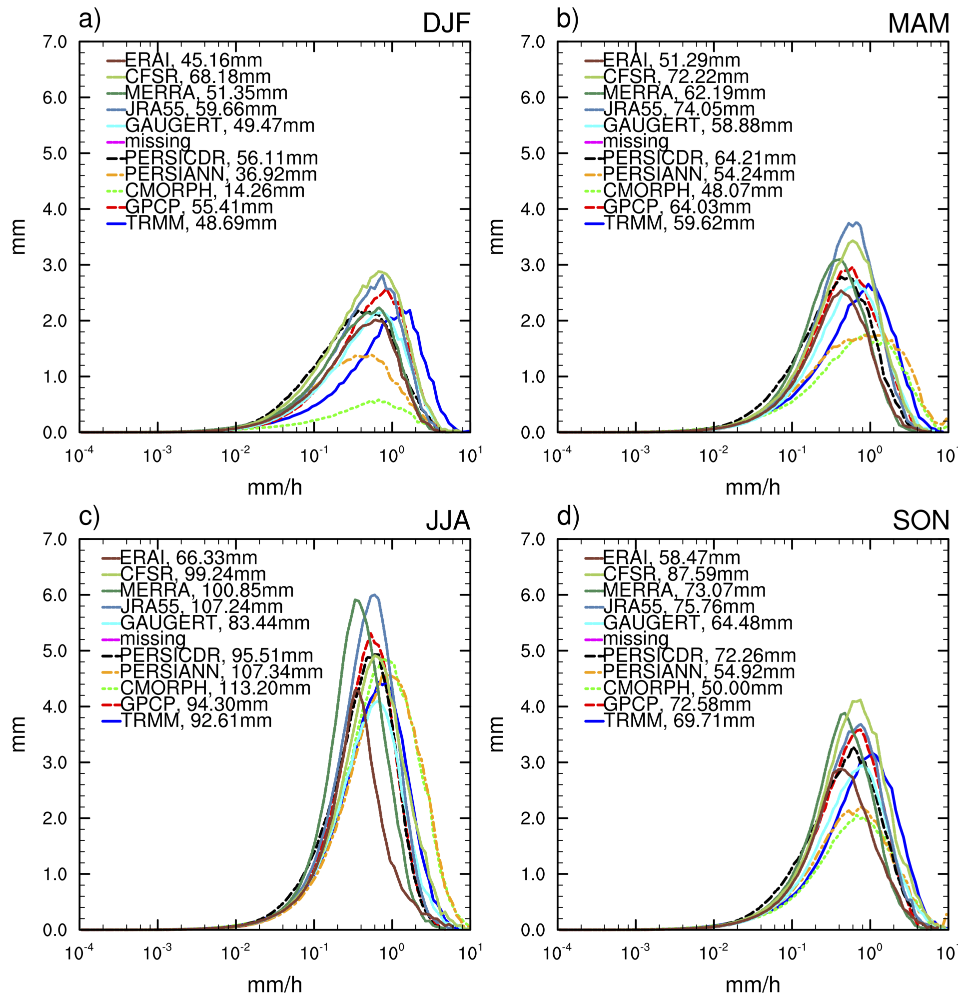

FIG. 12. Distribution of precipitation amount by precipitation rate over land area for North America (15°N - 49°N, the same area as is used in Fig. **??**). Panels a)-d) show the precipitation amount distribution for all seasons for 2006 - 2012. The average is computed over the years 2006 - 2012. Insets show average monthly totals during each season for the different estimates.

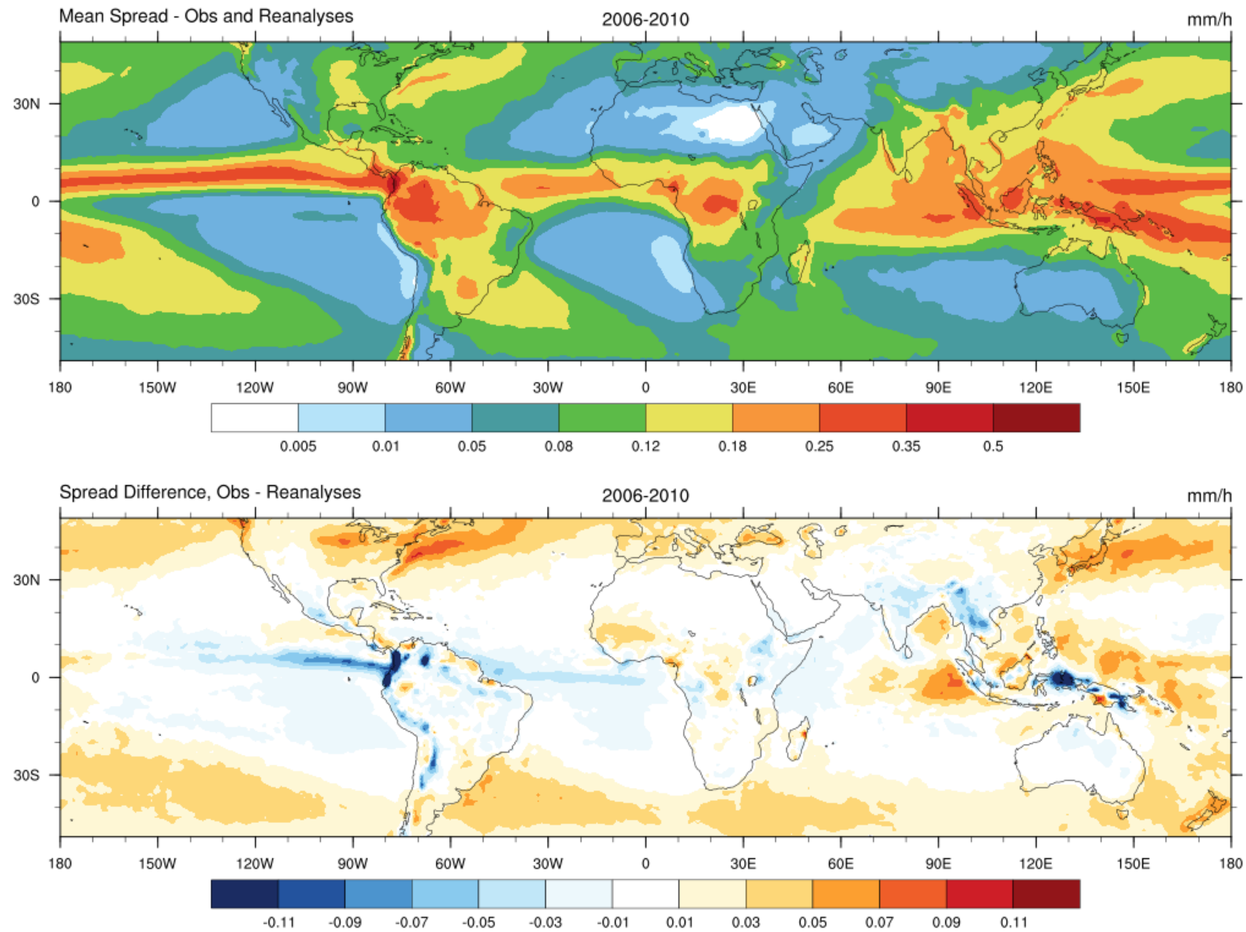

FIG. 13. Spread among precipitation estimates (computed as the mean standard deviation among data sets) for 2006-2010. Top panel: spread among precipitation data sets (including reanalyses). Bottom panel: difference in spread among observational precipitation data sets and spread among reanalyses. The mean seasonal cycle is removed from daily data prior to computing the spread.